

QUANTUM MONTE CARLO SIMULATIONS OF REAL SOLIDS

W. M. C. Foulkes¹, M. Nekovee¹, R. L. Gaudoin¹,
M. L. Stedman¹, R. J. Needs², R. Q. Hood²,
G. Rajagopal², M. D. Towler², P. R. C. Kent²,
Y. Lee², W.-K. Leung², A. R. Porter², and S. J. Breuer³

¹Blackett Laboratory, Imperial College

²Cavendish Laboratory, Cambridge University

³Edinburgh Parallel Computing Centre, University of Edinburgh

INTRODUCTION

Computer simulation has become a standard tool in almost all areas of science, engineering and technology, and its pervasive importance was one of the clearest themes to have emerged from the recent U.K. Technology Foresight Exercise. The Foresight Steering Group Report emphasised that the impressive recent advances in our understanding of complex systems are largely attributable to the use of computer simulation, and identified the field as a main priority for the future.

Although computer simulation is used in many different contexts, the underlying questions are often surprisingly similar. What happens when many simple objects come together and interact? How does the complex behaviour of the whole emerge from the simple laws obeyed by its parts? The constituent objects may range from the electrons and nuclei making up a crystal of silicon to the cables and girders holding up the Millennium Dome, but the common aim is to predict the complex large scale behaviour from the simpler small scale behaviour. In essence, computers are used to build bridges between different length scales.

Figure 1 shows a wide selection of different areas of physics, chemistry and engineering linked by arrows corresponding to changes of length scale. Computer simulations are already being used to help climb most of the arrows, and a large part of the work being discussed at this conference fits somewhere on the diagram. At the top level of the tree, the technological and economic importance of computer simulation is large and growing rapidly. Aircraft, cars, and bridges are already designed using computers, and new drugs and materials soon will be. Our work, being at the bottom of the tree, may appear far from the wealth creating activities at the top, but it provides the scientific understanding needed to underpin progress at higher levels.

To understand why work at the bottom of the tree is interesting and important, think about the problems facing someone higher up. Imagine, for example, that you're a materials scientist trying to use computer simulations to assess the effects of different

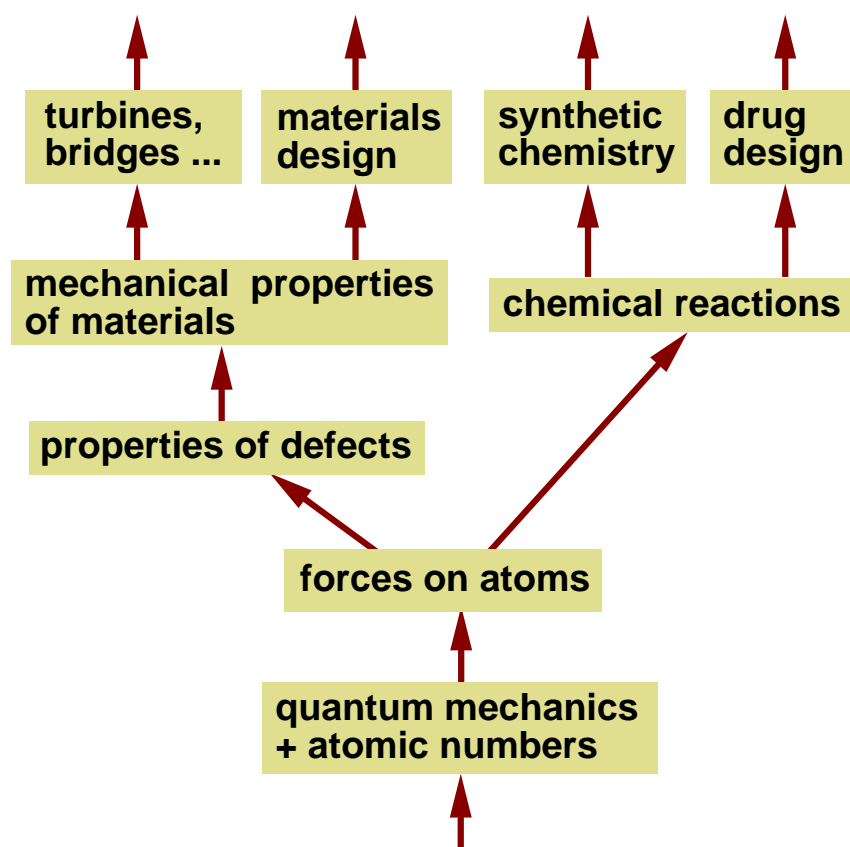


Figure 1. The simulation tree

manufacturing processes on the properties of steel. Since the mechanical behaviour of steel is governed by the dislocations, impurities, grain boundaries, and other defects introduced during manufacture, your job is to model a complicated system of interacting defects. Your starting point is a thorough quantitative understanding of the comparatively simple properties of single defects and interacting pairs of defects, but where is this understanding going to come from? Some of the information you need can probably be measured, but experiments are difficult, expensive and limited. The rest can only be obtained from the results of simulations one level lower on the tree. You can't begin to model steel without first modelling its constituent defects.

Almost everywhere you look on the tree, some of the data required for simulations at one level can only be obtained from the results of simulations on the level below. The descent continues right down to the quantum mechanical atomic/molecular level, where the many-electron Schrödinger equation at last provides a concise and accurate universal law of nature. There are quantum field theoretical levels below it, of course, but the Schrödinger equation requires no input from these and is accurate enough to explain almost everything higher up the tree. Since the quantum mechanical length scale at the bottom of Figure 1 is the largest at which we can write down a reliable “grand unified theory”, it serves as a natural root for the tree above it.

The “Consortium for Computational Quantum Many-Body Theory” was one of the most sharply focused of the consortia set up under the High Performance Computing Initiative (HPCI). Our goal is the development of a deeper conceptual understanding of the role of quantum mechanics in determining the electronic structure and properties of real materials. In pursuit of this goal, we're developing accurate and reliable methods for climbing from the lowest box in Figure 1 to the levels above. Our quantum me-

chemical calculations start from “first principles” and require only the atomic numbers of the atoms involved as input; if they live up to expectations, they should provide a firm base on which the higher levels of the simulation tree can rest.

Although the methods we use are very different, our goals are strikingly similar to those of the larger UKCP Consortium. In the longer term, we believe that the many-electron approach we have adopted¹ will supersede the one-electron approach² they favour, but this is a matter of conjecture of course. The good news is that both our Consortium and the UKCP Consortium are world leaders in their fields, and that the United Kingdom is well placed to seize the great opportunities offered by first principles materials simulation over the next twenty years. Within that time scale, we are confident that quantum mechanical simulations will be used to design drugs, catalysts, chemical syntheses, advanced materials, and electronic devices.

THE CHALLENGE

Our “grand unified theory”, the many-electron Schrödinger equation, is generally believed capable of predicting and describing almost everything experienced in everyday life. It doesn’t matter whether you’re interested in the biochemical reactions taking place in the human body, the flow of electrons through a transistor, or the strength of steel — they’re all in there somewhere.

Given its enormous multitude of complicated ramifications, the equation itself is surprisingly simple,

$$\left(-\frac{1}{2} \sum_{i=1}^N \nabla_i^2 + \sum_{i=1}^N v(\mathbf{r}_i) + \frac{1}{2} \sum_{i=1}^N \sum_{\substack{j=1 \\ (j \neq i)}}^N \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} \right) \Psi(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N) = E \Psi(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N) \quad .$$

It’s a linear second-order complex partial differential equation of a well-known type, the only significant computational problem being that the sums over i and j include terms for every electron in the system. An object weighing a gram or two contains something like 10^{23} electrons, so the many-electron Schrödinger equation is a partial differential equation in roughly 10^{23} variables. Most scientific and engineering problems involve partial differential equations in 3 or at most 6 variables, so dealing with 10^{23} is quite a challenge.

The traditional response to this challenge has been to stand back in horror and start approximating. Instead of considering all the electrons together, the idea has been to look at the electrons one by one, replacing the complicated fluctuating forces due to the others by an average force known as a mean field. This trick reduces the 10^{23} dimensional many-electron Schrödinger equation to a much simpler three-dimensional equation for each electron. The mean field depends on the electron density, which isn’t known until the one-electron equation has been solved to find all the one-electron quantum states, so it’s necessary to use some sort of iterative procedure to home in on a consistent solution.

The mean-field theory most often used in solids is density functional theory.² Although exact in principle, density functional theory involves an unknown quantity called the exchange-correlation energy functional which has to be approximated in practice. Several rival approximations are in widespread use, but the simplest and best known is the local density approximation (LDA). The LDA and its alternatives work surprisingly well in many cases, and density functional calculations are already being used to investigate questions of genuine industrial and economic importance. This is the province

of the UKCP Consortium, the members of which apply density functional theory to a wide range of topics in condensed matter physics, materials science, geophysics and biochemistry. Unfortunately, the LDA is by no means universally accurate or reliable. When it works, as in simple metals and semiconductors, it seems to work very well indeed; but when it doesn't, as in transition metal oxides, high temperature superconductors, and hydrogen-bonded materials such as water, there isn't much one can do. There's no simple systematic way to improve the standard density functional approach and one is forced to go back to the full many-electron Schrödinger equation instead.

QUANTUM MONTE CARLO

The main approaches used to solve the many-electron Schrödinger equation in real solids are quantum Monte Carlo¹ (QMC) and the *GW* method.³ The two are complementary — QMC simulations tell you about ground state properties while the *GW* method gives information about excitations — and we use both. In this article we'll concentrate on our QMC work since it's slightly easier to explain and particularly straightforward to implement on massively parallel computers. The generic term “quantum Monte Carlo” covers several different methods, but the ones we use are called variational QMC and diffusion QMC.

Variational QMC

According to quantum mechanics, the probability that a measurement of the positions of all N electrons in a solid finds them at $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N$ is proportional to $|\Psi(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)|^2$, where $\Psi(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ is the many-electron wavefunction. The idea behind variational QMC is to use a computer to generate sets of random positions distributed in exactly the same way as the results of this idealised measurement, and to average the outcomes of many such computer experiments to obtain quantum mechanical expectation values. Given the form of the many-electron wavefunction one can generate the required samples using the well-known Metropolis algorithm, but the exact many-electron wavefunction is an unknown function of an enormously large number of variables and has to be approximated.

The first approximation is to replace the macroscopic piece of solid containing roughly 10^{23} electrons by a small model system containing no more than a few thousand. This sounds drastic, but in fact it works very well, especially when we apply periodic boundary conditions (which means that an electron leaving one face of the model system immediately re-enters through the opposite face) to get rid of any surfaces. The replacement of an effectively infinite solid by a small model system subject to periodic boundary conditions is still an approximation, of course, and the associated finite-size errors caused us significant problems at the beginning. Luckily, however, our attempts to understand these errors have been very successful⁴ and we have been able to reduce them to the point that they can safely be ignored for most purposes. We consider this an important advance and expect the techniques we have developed to become standard practice in the field.

The second approximation is more problematic and harder to improve. Since we don't know the exact many-electron wavefunction (even for our small model system) we have to guess it. At first glance this looks like a hopeless task, but a surprisingly large fraction of those few quantum many-body problems that have ever been “solved” have in fact been solved by guessing the wavefunction (think of the BCS theory of superconductivity and Laughlin's theory of the fractional quantum Hall effect). Most

of our calculations are done using trial functions of the Slater-Jastrow type,

$$\Psi_T = D^\uparrow D^\downarrow \exp \left[- \sum_{i>j} u(\mathbf{r}_i, \mathbf{r}_j) + \sum_i \chi(\mathbf{r}_i) \right] ,$$

where D^\uparrow and D^\downarrow are Slater determinants of spin-up and spin-down single-particle orbitals obtained from Hartree-Fock or LDA calculations, the function $u(\mathbf{r}_i, \mathbf{r}_j)$ correlates the motion of pairs of electrons, and $\chi(\mathbf{r}_i)$ is a one-body function as introduced by Fahy.⁵ The Slater determinants build in the antisymmetry required by the Pauli principle, and the u and χ functions are adjusted to minimise the total energy (or more precisely⁶ the variance of the total energy) in accordance with the variational principle.

This simple trial wavefunction, first applied to real solids at the end of the eighties,⁵ has proved astonishingly accurate for the weakly correlated solids studied so far. Following careful optimisation of the u and χ functions, total and cohesive energies are accurate to within 0.2eV per atom, approximately five times better than the best LDA calculations.⁷ The full range of solids for which this trial wavefunction is accurate is still not known, but so far it has exceeded all reasonable expectations.

Diffusion QMC

The major limitation of the variational QMC method is obvious: what happens when the assumed trial wavefunction isn't accurate enough? The direct approach is to add more variational parameters and resort to brute force optimisation, but this limits one to an assumed functional form which may not be adequate, particularly in strongly correlated systems where the quantum state is markedly different from the Fermi liquid state found in weakly correlated solids.

An alternative and much better approach is diffusion QMC, which is based on the imaginary-time Schrödinger equation,

$$\left(-\frac{1}{2} \nabla_{\mathbf{R}}^2 + V(\mathbf{R}) \right) \Psi(\mathbf{R}, \tau) = -\frac{\partial \Psi(\mathbf{R}, \tau)}{\partial \tau} ,$$

where we have adopted a very condensed notation in which $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ is a $3N$ -dimensional vector containing all the electron positions, $\nabla_{\mathbf{R}}^2$ is shorthand for $\sum_i \nabla_i^2$, and $V(\mathbf{R})$ is the sum of all the potential energy terms appearing in the ordinary many-electron Schrödinger equation. The imaginary-time Schrödinger equation reduces to the usual time-dependent Schrödinger equation when τ is replaced by the imaginary variable it , but we will be interested only in real τ from now on.

It's straightforward to show that as long as the starting state, $\Psi(\mathbf{R}, \tau = 0)$, isn't orthogonal to the ground state energy eigenfunction, the solution of the imaginary-time Schrödinger equation becomes proportional to the ground state in the limit as $\tau \rightarrow \infty$. The imaginary-time development is just a mathematical trick used to convert an arbitrary starting state into the ground state without assuming any particular functional form.

The clever part is that there's a simple algorithm that can be used to carry out the imaginary-time development in systems containing hundreds or thousands of electrons. If the $V(\mathbf{R})$ term is ignored, the imaginary-time Schrödinger equation is just a diffusion equation describing a population of "walkers" diffusing in a $3N$ -dimensional "configuration space"; if the $\nabla_{\mathbf{R}}^2$ term is ignored, it's just a rate equation analogous to the equation describing radioactive decay. Putting these two things together, it can be seen that the imaginary-time Schrödinger equation describes walkers diffusing in $3N$ -dimensional space while dying out or multiplying at a rate determined by the

value of $V(\mathbf{R})$. This simple physical picture leads directly to a simulation method in which a collection of walkers multiply and reproduce as they diffuse randomly around in a $3N$ -dimensional vector space. After a long time τ , the density of diffusing walkers becomes proportional to the ground state wavefunction $\Psi_0(\mathbf{R})$ (not $|\Psi_0(\mathbf{R})|^2$ as you might expect; it's $\Psi_0(\mathbf{R})$ that plays the role of a probability density here), and one can start accumulating ground state expectation values as in variational QMC.

There's one major problem with the algorithm just described. The diffusion QMC method finds the overall ground state of the system, and for the many-electron Schrödinger equation this happens to be a totally symmetric function of the particle coordinates. Electrons are fermions, however, and according to the Pauli principle many-fermion states have to be totally antisymmetric. The simple diffusion QMC algorithm gives a many-boson state of no physical interest.

To ensure that the diffusion QMC simulation produces an antisymmetric state we're forced to make the so-called fixed node approximation, in which we guess the shapes of the regions of configuration space within which the wavefunction is positive and negative and solve the imaginary-time Schrödinger equation to find the lowest energy state consistent with that guess. Fixed-node diffusion QMC is perhaps best regarded as a variational method in which, instead of assuming a trial form for the whole wavefunction, we assume only a trial form for the nodal surface (the $3N - 1$ dimensional configuration-space surface on which the wavefunction is zero). We use fairly simple guesses based on mean-field wavefunctions and our experience so far suggests that these give very good results, accurate to considerably better than 0.1eV per atom.⁷ The fixed-node approximation is the only significant approximation in our calculations.

COMPUTATIONAL ISSUES

Like most Monte Carlo methods, our QMC calculations are ideally suited to massively parallel architectures. We simply put one walker (or several walkers) on each node and run different simulations on each. Some communication between the many parallel simulations is necessary, but the communications requirements are quite small and the performance increases linearly with the number of nodes used.

One of the comments we received in response to our first HPCI grant proposal was that our work was so “embarrassingly parallel” that it wasn't worth running on the Cray T3D; it wouldn't provide enough of a technical challenge! In our view this is nonsense. Our use of the T3D and T3E has allowed us to do world-beating physics which would not have been possible on lesser machines or networks of workstations. Yes, our algorithms are easy to parallelise, but surely this is a strength rather than a weakness? It means that we can get our programs running quickly and efficiently and concentrate our attention on the science instead of on its computational implementation. The T3D and T3E have enabled us to study larger systems and achieve higher accuracies than any other group in the world, and to complete ambitious scientific projects that nobody else could have attempted. Our successes have demonstrated the tremendous potential of the combination of quantum Monte Carlo methods and massively parallel computation. This, we believe, is the sort of work the T3D and T3E were intended for.

We have, of course, put considerable effort into optimising our codes, and thanks to the expert help we have received from Edinburgh Parallel Computing Centre (EPCC) they are now amongst the most efficient on the T3D and T3E. We have found it very useful to have a named EPCC postdoc (currently Stephen Breuer) assigned to look after our Consortium, and hope that the EPSRC will provide funds to continue this system in future.

SOME EXAMPLE APPLICATIONS

Our consortium has been involved in a wide variety of different projects, leading to 25 publications in refereed international journals, 6 of which were in Physical Review Letters. Our access to the T3D and T3E has allowed us to establish a position as one of the world’s two leading groups in our field, and we have recently been invited to write major review articles in both “Reports on Progress in Physics” and “Advances in Chemical Physics”. We plan to use these to explain the techniques we have developed in the hope that this will help to demystify the subject and encourage its further growth. We can’t possibly cover all our QMC work in this short paper, so we’ve picked four representative examples.

Total Energy Calculations

Our first few projects were concerned with testing the accuracy and reliability of QMC calculations of ground state energies.⁷ We used QMC to calculate the cohesive energies of a number of different solids and compared the results against mean-field values obtained using density functional theory within the LDA. Typical results are summarised in Table 1.

Table 1. The Cohesive Energy of Ge Obtained Using Three Different Methods

Method Used	Cohesive Energy (eV/atom)
LDA Calculation	4.59
Diffusion QMC Calculation	3.85
Experiment	3.85

Finite-size errors in QMC simulations are typically a significant fraction of an eV, so highly accurate calculations such as that in Table 1 rely on the techniques we have developed for eliminating these errors;⁴ it’s only after the finite-size errors have been eliminated that the full accuracy of the QMC approach becomes apparent. For weakly correlated solids like Ge, it’s now clear that diffusion QMC calculations of total energy differences (and hence interatomic forces) are at least an order of magnitude more accurate than calculations using any alternative method.

Relativistic Electron Gas

Close to the nuclei of heavy atoms, where the electron density is very high, the electrons move at a significant fraction of the speed of light and relativistic effects become important. These are not included in the many-electron Schrödinger equation, but can be incorporated via perturbation theory. We have been studying the relativistic effects in various systems to order $1/c^2$ using QMC methods. After a starter project in which we studied the relativistic effects in various atoms,⁸ we concentrated on the unpolarised relativistic electron gas. Our results enabled us to check the relativistic version of the LDA used in density functional calculations of solids and molecules containing heavy atoms. We found that the contact term is well described by the LDA, but that the effects of retardation are very poorly described and require a non-local functional. A short description of this work⁹ appeared in Physical Review Letters, and we are currently writing up our new results on the polarised relativistic homogeneous electron gas.

Silicon Quasiparticle Bandstructure

This project demonstrated the power of diffusion QMC by calculating the band structure of silicon, obtaining excellent agreement with experiment.¹⁰ Although similar calculations had been done in the past,¹¹ ours were much more complete and precise than their predecessors and were the first to demonstrate the very high accuracy of the QMC approach. The excited state issues that can be addressed using QMC are somewhat limited, but in those cases where QMC can be applied it works very well and gives a good account of the many-body effects missed by mean-field methods. Unlike perturbative methods for studying excited states, QMC can also be used to study strongly correlated systems, and hence provides a unified framework for studying ground states and excitations throughout the periodic table. No other approach shows such promise.

Another result of this work was a better understanding of the theory underlying QMC calculations of excitation energies. A widely believed “folk theorem” states that the fixed-node diffusion QMC energy is always greater than or equal to the energy of the ground state eigenfunction with the same symmetry as the trial state used to define the nodal surface. In fact, we have recently proved that this theorem is correct only when the relevant ground state eigenfunction is non-degenerate, and have devised a simple analytic example to demonstrate our point.

Exact Density Functional Theory

In order to understand the physics behind mean-field approximations such as the LDA, it helps to think about one particular electron (let’s call it the labelled electron) moving through the sea of nuclei and other electrons making up a solid. The labelled electron is attracted to the positively charged nuclei, which are so massive that they can be treated as immobile, and repelled by the other electrons, which are negatively charged.

As a rough approximation, it seems sensible to replace the fluctuating forces due to the other electrons by the static electrical (Coulomb) force due to the *average* electronic charge density. This simple mean-field approximation, known as the Hartree approximation, helps keep the labelled electron away from regions where there are lots of other electrons on average, which is a good start, but misses something important. As the labelled electron moves around, the others stay out of its way; you can think of the labelled electron carrying round a little “exclusion zone”, usually known as the exchange-correlation (XC) hole, within which other electrons rarely venture. The electron density near the labelled electron is therefore less than the average density, and the Hartree approximation doesn’t take this into account. Other mean-field approaches such as Hartree-Fock theory attempt to build in the effects of the XC hole in an approximate way, but these approximations aren’t particularly accurate.

Density functional theory is based on a remarkable theorem, first proved almost 35 years ago,¹² which states that it is in principle possible to devise an exact mean-field theory. In other words, the mean field can be chosen in such a way that the energies and electron densities obtained by solving the one-electron equations come out exactly right. There’s even a prescription for constructing the exact mean field given the detailed shape of the XC hole.¹³ Practical applications of density functional theory have had to rely on approximations such as the LDA only because no available method has been able to calculate the shapes of XC holes in real materials.

Quantum Monte Carlo methods can calculate accurate XC holes, however, and Figure 2 shows the result of our calculation of the shape of the hole around a labelled




Figure 2. The exchange-correlation hole around an electron near the centre of a bond in crystalline silicon

electron at a particular point in the (110) plane of silicon.¹⁴ We’ve now assembled enough data to allow us to plot the shape of the XC hole anywhere in a piece of silicon, and have made a movie to show how the hole changes as the labelled electron moves around. The *exact* mean fields calculated from our XC hole data for silicon and other systems¹⁵ are significantly different from the commonly used approximations, and we believe that these results will help the development of better mean-field methods in the future.

CONCLUSIONS

This article was intended to emphasise the importance of first principles materials simulation and to explain where QMC fits in. We hope we have convinced you that the impact of quantum mechanical materials simulation will continue to grow in the future, and that QMC will be the method of choice when high accuracy is required.

Mean-field methods such as density functional theory within the LDA are fine when the necessary accuracy is a few tenths of an eV per atom, but many problems require much higher accuracy than this. Chemical reaction rates, for example, are sensitive to energy differences of order $k_B T \simeq 0.025$ eV at room temperature, and many biochemical processes require even higher accuracy. There is no doubt that quantum mechanical materials simulation promises a technical revolution, with computers used to design drugs, chemical syntheses, and smart materials, but this will not happen until accuracies of order 0.01 eV per atom can be attained routinely. No method applicable to solids or large molecules can yet provide such precision, but quantum Monte Carlo

is the closest and (in our opinion) the most likely to succeed.

Fortunately for us, it seems that the EPSRC concurs with this opinion. Computational quantum many-body theory, including both QMC calculations and methods for excited states such as the *GW* self-energy calculations that form the other main part of our work, were highlighted in the EPSRC “Review of Condensed Matter Physics” under “Future Growth Areas/Priorities”, and were featured in the “Analysis of Community Questionnaire” as a “likely major growth area in CMP research over the next 5 years”. The United Kingdom is among the world leaders in most areas of first principles materials simulation, and with the continuing support of the EPSRC there is no reason why we should not remain at the forefront.

REFERENCES

1. See, e.g., B. L. Hammond, W. A. Lester, and P. J. Reynolds. “Monte Carlo Methods in Ab Initio Quantum Chemistry,” World Scientific, Singapore, (1994); D. M. Ceperley and M. H. Kalos, *in*: “Monte Carlo Methods in Statistical Physics,” K. Binder, ed., Springer-Verlag, Berlin (1984); K. E. Schmidt and M. H. Kalos, *in*: “Applications of the Monte Carlo Method in Statistical Physics,” K. Binder, ed., Springer-Verlag, Berlin (1984); K. E. Schmidt, *in*: “Models and Methods in Few-Body Physics, Proceedings of the 8th Autumn School, Lisbon, Portugal, 13-18 Oct. 1986,” Springer-Verlag, Berlin (1987).
2. See, e.g., R. G. Parr and W. Yang. “Density-Functional Theory of Atoms and Molecules,” Oxford University Press, New York (1989); R. M. Dreizler and E. K. U. Gross. “Density Functional Theory, An Approach to the Quantum Many-Body Problem,” Springer-Verlag, Berlin (1990).
3. See, e.g., M. S. Hybertsen and S. G. Louie, *Phys. Rev. B.* 34:5390 (1986); R. W. Godby, M. Schluter, and L. J. Sham, *Phys. Rev. B.* 56:10159 (1988); L. Hedin and S. Lundquist, *Solid State Physics, Advances in Research and Applications* 23:1 (1969).
4. L. M. Fraser, W. M. C. Foulkes, G. Rajagopal, R. J. Needs, S. D. Kenny, and A. J. Williamson, *Phys. Rev. B.* 53:1814 (1996); A. J. Williamson, G. Rajagopal, R. J. Needs, L. M. Fraser, W. M. C. Foulkes, Y. Wang, and M.-Y. Chou, *Phys. Rev. B.* 55:4851 (1997).
5. S. Fahy, X. W. Wang, and S. G. Louie, *Phys. Rev. Lett.* 61:1631 (1988).
6. A. J. Williamson, S. D. Kenny, G. Rajagopal, A. J. James, R. J. Needs, L. M. Fraser, W. M. C. Foulkes, and P. Maccallum, *Phys. Rev. B.* 53:9640 (1996).
7. G. Rajagopal, R. J. Needs, S. Kenny, W. M. C. Foulkes, and A. James, *Phys. Rev. Lett.* 73:1959 (1994); G. Rajagopal, R. J. Needs, A. James, S. D. Kenny, and W. M. C. Foulkes, *Phys. Rev. B.* 51:10591 (1995).
8. S. D. Kenny, G. Rajagopal, and R. J. Needs, *Phys. Rev. A.* 51:1898 (1995).
9. S. D. Kenny, G. Rajagopal, R. J. Needs, W.-K. Leung, M. J. Godfrey, A. J. Williamson, and W. M. C. Foulkes, *Phys. Rev. Lett.* 77:1099 (1996).
10. A. J. Williamson, R. Q. Hood, R. J. Needs, and G. Rajagopal, accepted for publication in *Phys. Rev. B.* (1998); P. R. C. Kent, R. Q. Hood, M. D. Towler, R. J. Needs, and G. Rajagopal, submitted to *Phys. Rev. B.* (1997).
11. L. Mitás and R. M. Martin, *Phys. Rev. Lett.* 72:2438 (1994).
12. P. Hohenberg and W. Kohn, *Phys. Rev.* 136:B864 (1964); W. Kohn and L. J. Sham, *Phys. Rev.* 140:A1133 (1965).
13. J. Harris and R. O. Jones, *J. Phys. F.* 4:1170 (1974); D. C. Langreth and J. P. Perdew, *Phys. Rev. B.* 21:5469 (1980); O. Gunnarson and B. I. Lundqvist, *Phys. Rev. B.* 13:4274 (1976).
14. R. Q. Hood, M.-Y. Chou, A. J. Williamson, G. Rajagopal, R. J. Needs, and W. M. C. Foulkes, *Phys. Rev. Lett.* 78:3350 (1997); R. Q. Hood, M.-Y. Chou, A. J. Williamson, G. Rajagopal, and R. J. Needs, accepted for publication in *Phys. Rev. B.* (1998).
15. M. Nekovee, W. M. C. Foulkes, A. J. Williamson, G. Rajagopal, and R. J. Needs, accepted for publication in *Advances in Quantum Chemistry* (1997).